
1 Data Augmentation via Lévy Processes

Stefan Wager

Stanford University
Stanford, USA

swager@stanford.edu

William Fithian

University of California, Berkeley
Berkeley, USA

wfithian@berkeley.edu

Percy Liang

Stanford University
Stanford, USA

pliang@cs.stanford.edu

If a document is about travel, we may expect that short snippets of the document should also be about travel. We introduce a general framework for incorporating these types of invariances into a discriminative classifier. The framework imagines data as being drawn from a slice of a Lévy process. If we slice the Lévy process at an earlier point in time, we obtain additional pseudo-examples, which can be used to train the classifier. We show that this scheme has two desirable properties: it preserves the Bayes decision boundary, and it is equivalent to fitting a generative model in the limit where we rewind time back to 0. Our construction captures popular schemes such as Gaussian feature noising and dropout training, as well as admitting new generalizations.

1.1 Introduction

Black-box discriminative classifiers such as logistic regression, neural networks, and SVMs are the go-to solution in machine learning: they are simple to apply and often perform well. However, an expert may have additional knowledge to exploit, often taking the form of a certain family of transformations that should usually leave labels fixed. For example, in object recognition, an image of a cat rotated, translated, and peppered with a small amount of noise is probably still a cat. Likewise, in document classification, the first paragraph of an article about travel



(a) Gaussian noise

The colorful Norwegian city of Bergen is also a gateway to majestic fjords. Bryggen Hanseatic Wharf will give you a sense of the local culture – take some time to snap photos of the Hanseatic commercial buildings, which look like scenery from a movie set.



The colorful of gateway to fjords. Hanseatic Wharf will sense the culture – take some to snap photos the commercial buildings, which look scenery a

(b) Dropout noise

Figure 1.1: Two examples of transforming an original input X into a noisy, less informative input \tilde{X} . The new inputs clearly have the same label but contain less information and thus are harder to classify.

is most likely still about travel. In both cases, the “expert knowledge” amounts to a belief that a certain transform of the features should generally not affect an example’s label.

One popular strategy for encoding such a belief is *data augmentation*: generating additional pseudo-examples or “hints” by applying label-invariant transformations to training examples’ features (Abu-Mostafa, 1990; Schölkopf et al., 1997; Simard et al., 1998). That is, each example $(X^{(i)}, Y^{(i)})$ is replaced by many pairs $(\tilde{X}^{(i,b)}, Y^{(i)})$ for $b = 1, \dots, B$, where each $\tilde{X}^{(i,b)}$ is a transformed version of $X^{(i)}$. This strategy is simple and modular: after generating the pseudo-examples, we can simply apply any supervised learning algorithm to the augmented dataset. Figure 1.1 illustrates two examples of this approach, an image transformed to a noisy image and a text caption, transformed by deleting words.

Dropout training (Srivastava et al., 2014) is an instance of data augmentation that, when applied to an input feature vector, zeros out a subset of the features randomly. Intuitively, dropout implies a certain amount of signal redundancy across features—that an input with about half the features masked should usually be classified the same way as a fully-observed input. In the setting of document clas-

sification, dropout can be seen as creating pseudo-examples by randomly omitting some information (i.e., words) from each training example. Building on this interpretation, Wager et al. (2014) show that learning with such artificially difficult examples can substantially improve the generalization performance of a classifier.

To study dropout, Wager et al. (2014) assume that documents can be summarized as Poisson word counts. Specifically, assume that each document has an underlying topic associated with a word frequency distribution π on the p -dimensional simplex and an expected length $T \geq 0$, and that, given π and T , the word counts X_j are independently generated as $X_j | T, \pi \sim \text{Pois}(T\pi_j)$. The analysis of Wager et al. (2014) then builds on a duality between dropout and the above generative model. Consider the example given in Figure 1.1, where dropout creates pseudo-documents \tilde{X} by deleting half the words at random from the original document X . As explained in detail in Section 1.2.1, if X itself is drawn from the above Poisson model, then the dropout pseudo-examples \tilde{X} are marginally distributed as $\tilde{X}_j | T, \pi \sim \text{Pois}(0.5T\pi_j)$. Thus, in the context of this Poisson generative model, dropout enables us to create new, shorter pseudo-examples that preserve the generative structure of the problem.

The above interpretation of dropout raises the following question: if feature deletion is a natural way to create information-poor pseudo-examples for document classification, are there natural analogous feature noising schemes that can be applied to other problems? In this chapter, we seek to address this question, and study a more general family of data augmentation methods generalizing dropout, based on Lévy processes: We propose an abstract Lévy thinning scheme that reduces to dropout in the Poisson generative model considered by Wager et al. (2014). Our framework further suggests new methods for feature noising such as Gamma noising based on alternative generative models, all while allowing for a unified theoretical analysis.

From generative modeling to data augmentation. In the above discussion, we treated the expected document length T as fixed. More generally, we could imagine the document as growing in length over time, with the observed document X merely a “snapshot” of what the document looks like at time T . Formally, we can imagine a latent Poisson process $(A_t)_{t \geq 0}$, with fixed- t marginals $(A_t)_j | \pi \sim \text{Pois}(t\pi_j)$, and set $X = A_T$. In this notation, dropout amounts to “rewinding” the process A_t to obtain short pseudo-examples. By setting $\tilde{X} = A_{\alpha T}$, we have $\mathbb{P}[\tilde{X} = \tilde{x} | X = x] = \mathbb{P}[A_{\alpha T} = \tilde{x} | A_T = x]$, for thinning parameter $\alpha \in (0, 1)$.

The main result of this chapter is that the analytic tools developed by Wager et al. (2014) are not restricted to the case where (A_t) is a Poisson process, and in fact hold whenever (A_t) is a *Lévy process*. In other words, their analysis applies to any classification problem where the features X can be understood as time- T snapshots of a process (A_t) , i.e., $X = A_T$.

Recall that a Lévy process $(A_t)_{t \geq 0}$ is a stochastic process with $A_0 = 0$ that has independent and stationary increments: $\{A_{t_i} - A_{t_{i-1}}\}$ are independent for $0 = t_0 < t_1 < t_2 < \dots$, and $A_t - A_s \stackrel{d}{=} A_{t-s}$ for and $s < t$. Common examples

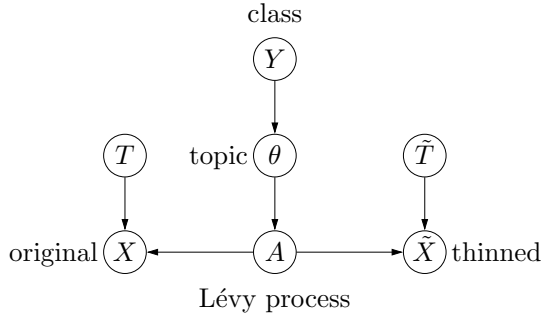


Figure 1.2: Graphical model depicting our generative assumptions; note that we are not fitting this generative model. Given class Y , we draw a topic θ , which governs the parameters of the Lévy process (A_t) . We slice at time T to get the original input $X = A_T$ and at an earlier time \tilde{T} to get the thinned or noised input $\tilde{X} = A_{\tilde{T}}$. We show that given X , we can sample \tilde{X} without knowledge of θ .

of Lévy processes include Brownian motion and Poisson processes.

In any such Lévy setup, we show that it is possible to devise an analogue to dropout that creates pseudo-examples by *rewinding* the process back to some earlier time $\tilde{T} \leq T$. Our generative model is depicted in Figure 1.2: (A_t) , the information relevant to classifying Y , is governed by a latent topic $\theta \in \mathbb{R}^p$. Lévy thinning then seeks to rewind (A_t) —importantly as we shall see, without having access to θ .

We should think of (A_t) as representing an ever-accumulating amount of information concerning the topic θ : In the case of document classification, (A_t) are the word counts associated with a document that grows longer as t increases. In other examples that we discuss in Section 1.3, A_t will represent the sum of t independent noisy sensor readings. The independence of increments property assures that as we progress in time, we are always obtaining new information. The stopping time T thus represents the *information content* in input X about topic θ . Lévy thinning seeks to improve classification accuracy by turning a few information-rich examples X into many information-poor examples \tilde{X} .

We emphasize that, although our approach uses generative modeling to motivate a data augmentation scheme, we do not in fact fit a generative model. This presents a contrast to the prevailing practice: two classical approaches to multiclass classification are to either directly train a discriminative model by running, e.g., multiclass logistic regression on the n original training examples; or, at the other extreme, to specify and fit a simple parametric version of the above generative model, e.g., naive Bayes, and then use Bayes’ rule for classification. It is well known that the latter approach is usually more efficient if it has access to a correctly specified generative model, but may be badly biased in case of model misspecification (Efron, 1975; Ng and Jordan, 2002; Liang and Jordan, 2008). Here, we first seek to devise a noising scheme $X \rightarrow \tilde{X}$ and then to train a discriminative model on the pseudo-examples (\tilde{X}, Y) instead of the original examples (X, Y) . Note




Figure 1.3: We model each input X as a slice of a Lévy process at time T . We generate noised examples \tilde{X} by “stepping back in time” to \tilde{T} . Note that the examples of the two classes are closer together now, thus forcing the classifier to work harder.

that even if the generative model is incorrect, this approach will incur limited bias as long as the noising scheme roughly preserves class boundaries — for example, even if the Poisson document model is incorrect, we may still be justified in classifying a subsampled travel document as a travel document. As a result, this approach can take advantage of an abstract generative structure while remaining more robust to model misspecification than parametric generative modeling.

Overview of results. We consider the multiclass classification setting where we seek to estimate a mapping from input X to class label Y . We imagine that each X is generated by a mixture of Lévy process, where we first draw a random topic θ given the class Y , and then run a Lévy process (A_t) depending on θ to time T . In order to train a classifier, we pick a thinning parameter $\alpha \in (0, 1)$, and then create pseudo examples by rewinding the original X back to time αT , i.e., $\tilde{X} \sim A_{\alpha T} \mid A_T$.

We show three main results in this chapter. Our first result is that we can generate such pseudo-examples \tilde{X} without knowledge of the parameters θ governing the generative Lévy process. In other words, while our method posits the existence of a generative model, our algorithm does not actually need to estimate it. Instead, it enables us to give hints about a potentially complex generative structure to a discriminative model such as logistic regression.

Second, under assumptions that our generative model is correct, we show that

feature noising preserves the Bayes decision boundary: $\mathbb{P}[Y \mid X = x] = \mathbb{P}[Y \mid \tilde{X} = x]$. This means that feature noising does not introduce any bias in the infinite data limit.

Third, we consider the limit of rewinding to the beginning of time ($\alpha \rightarrow 0$). Here, we establish conditions given which, even with finite data, the decision boundary obtained by fitting a linear classifier on the pseudo-examples is equivalent to that induced by a simplified generative model. When this latter result holds, we can interpret α -thinning as providing a semi-generative regularization path for logistic regression, with a simple generative procedure at one end and unregularized logistic regression at the other.

Related work. The trade-off between generative models and discriminative models has been explored extensively. Rubinstein and Hastie (1997) empirically compare discriminative and generative classifiers models with respect to bias and variance, Efron (1975) and Ng and Jordan (2002) provide a more formal discussion of the bias-variance trade-off between logistic regression and naive Bayes. Liang and Jordan (2008) perform an asymptotic analysis for general exponential families.

A number of papers study hybrid loss functions that combine both a joint and conditional likelihood (Raina et al., 2004; Bouchard and Triggs, 2004; Lasserre et al., 2006; McCallum et al., 2006; Liang and Jordan, 2008). The data augmentation approach we advocate in this chapter is fundamentally different, in that we are merely using the structural assumptions implied by the generative models to generate more data, and are not explicitly fitting a full generative model.

The present work was initially motivated by understanding dropout training (Srivastava et al., 2014), which was introduced in the context of regularizing deep neural networks, and has had much empirical success (Ba and Frey, 2013; Goodfellow et al., 2013; Krizhevsky et al., 2012; Wan et al., 2013). Many of the regularization benefits of dropout can be found in logistic regression and other single-layer models, where it is also known as “blankout noise” (Globerson and Roweis, 2006; van der Maaten et al., 2013) and has been successful in natural language tasks such as document classification and named entity recognition (Wager et al., 2013; Wang and Manning, 2013; Wang et al., 2013). There are a number of theoretical analyses of dropout: using PAC-Bayes framework (McAllester, 2013), comparing dropout to “altitude training” (Wager et al., 2014), and interpreting dropout as a form of adaptive regularization (Baldi and Sadowski, 2014; Bishop, 1995; Helmbold and Long, 2015; Josse and Wager, 2014; Wager et al., 2013).

1.2 Lévy Thinning

We begin by briefly reviewing the results of Wager et al. (2014), who study dropout training for document classification from the perspective of thinning documents (Section 1.2.1). Then, in Section 1.2.2, we generalize these results to the setting of

generic Lévy generative models.

1.2.1 Motivating Example: Thinning Poisson Documents

Suppose we want to classify documents according to their subject, e.g., sports, politics, or travel. As discussed in the introduction, common sense intuition about the nature of documents suggests that a short snippet of a sports document should also be classified as a sports document. If so, we can generate many new training examples by cutting up the original documents in our dataset into shorter subdocuments and labeling each subdocument with the same label as the original document it came from. By training a classifier on all of the pseudo-examples we generate in this way, we should be able to obtain a better classifier.

In order to formalize this intuition, we can represent a document as a sequence of words from a dictionary $\{1, \dots, d\}$, with the word count X_j denoting the number of occurrences of word j in the document. Given this representation, we can easily create “subdocuments” by binomially downsampling the word counts X_j independently. That is, for some fixed downsampling fraction $\alpha \in (0, 1)$, we draw

$$\tilde{X}_j \mid X_j \sim \text{Binom}(X_j, \alpha). \quad (1.1)$$

In other words, we keep each occurrence of word j independently with probability α .

Wager et al. (2014) study this downsampling scheme in the context of a Poisson mixture model for the inputs X that obeys the structure of Figure 1.2: first, we draw a class $Y \in \{1, \dots, K\}$ (e.g., travel) and a “topic” $\theta \in \mathbb{R}^d$ (e.g., corresponding to travel in Norway). The topic θ specifies a distribution over words,

$$\mu_j(\theta) = e^{\theta_j}, \quad (1.2)$$

where, without loss of generality, we assume that $\sum_{j=1}^d e^{\theta_j} = 1$. We then draw a $\text{Pois}(T)$ number of words, where T is the expected document length, and generate each word independently according to θ . Equivalently, each word count is an independent Poisson random variable, $X_j \sim \text{Pois}(T\mu_j(\theta))$. The following is an example draw of a document:

$$\begin{aligned} Y &= \text{travel} \\ \theta &= \left[\overbrace{0.5}^{\text{norway}}, \overbrace{0.5}^{\text{fjord}}, \overbrace{1.2}^{\text{the}}, \overbrace{-2.7}^{\text{skyscraper}}, \dots \right] \\ X &= \left[\overbrace{2}^{\text{norway}}, \overbrace{1}^{\text{fjord}}, \overbrace{3}^{\text{the}}, \overbrace{0}^{\text{skyscraper}}, \dots \right] \\ \tilde{X} &= \left[\overbrace{1}^{\text{norway}}, \overbrace{0}^{\text{fjord}}, \overbrace{1}^{\text{the}}, \overbrace{0}^{\text{skyscraper}}, \dots \right] \end{aligned}$$

Let us now try to understand the downsampling scheme $\tilde{X} \mid X$ in the context of the

Poisson topic model over X . For each word j , recall that $\tilde{X}_j \mid X_j \sim \text{Binom}(X_j, \alpha)$. If we marginalize over X , then we have:

$$\tilde{X}_j \mid T, \theta \sim \text{Pois}(\alpha T \mu_j(\theta)). \quad (1.3)$$

As a result, the distribution of \tilde{X} is exactly the distribution of X if we replaced T with $\tilde{T} = \alpha T$.

We can understand this thinning by embedding the document X in a multivariate Poisson process $(A_t)_{t \geq 0}$, where the marginal distribution of $A_t \in \{0, 1, 2, \dots\}^d$ is defined to be the distribution over counts when the expected document length is t . Then, we can write

$$X = A_T, \quad \tilde{X} = A_{\tilde{T}}. \quad (1.4)$$

Thus, under the Poisson topic model, the binomial thinning procedure does not alter the structure of the problem other than by shifting the expected document length from T to \tilde{T} . Figure 1.4 illustrates one realization of Lévy thinning in the Poisson case with a three-word dictionary. Note that in this case we can sample $\tilde{X} = A_{\alpha T}$ given $X = A_T$ without knowledge of θ .

This perspective lies at the heart of the analysis in Wager et al. (2014), who show under the Poisson model that, when the overall document length $\|X\|_1$ is independent of the topic θ , thinning does not perturb the optimal decision boundary. Indeed, the conditional distribution over class labels is identical for the original features and the thinned features:

$$\mathbb{P}[Y \mid X = x] = \mathbb{P}[Y \mid \tilde{X} = x]. \quad (1.5)$$

This chapter extends the result to general Lévy processes (see Theorem 1.2).

This last result (1.5) may appear quite counterintuitive: for example, if A_{60} is more informative than A_{40} , how can it be that downsampling does not perturb the conditional class probabilities? Suppose x is a 40-word document ($\|x\|_1 = 40$). When $t = 60$, most of the documents will be longer than 40 words, and thus x will be less likely under $t = 60$ than under $t = 40$. However, (1.5) is about the distribution of Y *conditioned on a particular realization* x . The claim is that, having observed x , we obtain the same information about Y regardless of whether t , the expected document length, is 40 or 60.

1.2.2 Thinning Lévy Processes

The goal of this section is to extend the Poisson topic model from Section 1.2.1 and construct general thinning schemes with the invariance property of (1.5). We will see that Lévy processes provide a natural vehicle for such a generalization: The Poisson process used to generate documents is a specific Lévy process, and binomial sampling corresponds to “rewinding” the Lévy process back in time.

Consider the multiclass classification problem of predicting a discrete class $Y \in \{1, \dots, K\}$ given an input vector $X \in \mathbb{R}^d$. Let us assume that the joint distribution

over (X, Y) is governed by the following generative model:

1. Choose $Y \sim \text{Mult}(\pi)$, where π is on the K -dimensional simplex.
2. Draw a topic $\theta \mid Y$, representing a subpopulation of class Y .
3. Construct a Lévy process $(A_t)_{t \geq 0} \mid \theta$, where $A_t \in \mathbb{R}^d$ is a potential input vector at time t .
4. Observe the input vector $X = A_T$ at a fixed time T .

While the Lévy process imposes a fair amount of structure, we make no assumptions about the number of topics, which could be uncountably infinite, or about their distribution, which could be arbitrary. Of course, in such an unconstrained non-parametric setting, it would be extremely difficult to adequately fit the generative model. Therefore, we take a different tack: We will use the structure endowed by the Lévy process to generate pseudo-examples for consumption by a discriminative classifier. These pseudo-examples implicitly encode our generative assumptions.

The natural way to generate a pseudo-example (\tilde{X}, Y) is to “rewind” the Lévy process (A_t) backwards from time T (recall $X = A_T$) to an earlier time $\tilde{T} = \alpha T$ for some $\alpha \in (0, 1)$ and define the thinned input as $\tilde{X} = A_{\tilde{T}}$. In practice, (A_t) is unobserved, so we draw \tilde{X} conditioned on the original input $X = A_T$ and topic θ . In fact, we can draw many realizations of $\tilde{X} \mid X, \theta$.

Our hope is that a single full example (X, Y) is rich enough to generate many different pseudo-examples (\tilde{X}, Y) , thus increasing the effective sample size. Moreover, Wager et al. (2014) show that training with such pseudo-examples can also lead to a somewhat surprising “altitude training” phenomenon whereby thinning yields an improvement in generalization performance because the pseudo-examples are more difficult to classify than the original examples, and thus force the learning algorithm to work harder and learn a more robust model.

A technical difficulty is that generating $\tilde{X} \mid X, \theta$ seemingly requires knowledge of the topic θ driving the underlying Lévy process (A_t) . In order to get around this issue, we establish the following condition under which the observed input $X = A_T$ alone is sufficient—that is, $\mathbb{P}[\tilde{X} \mid X, \theta]$ does not actually depend on θ .

Assumption 1.1 (exponential family structure). *The Lévy process $(A_t) \mid \theta$ is drawn according to an exponential family model whose marginal density at time t is*

$$f_{\theta}^{(t)}(x) = \exp[\theta \cdot x - t\psi(\theta)] h^{(t)}(x) \text{ for every } t \in \mathbb{R}. \quad (1.6)$$

Here, the topic $\theta \in \mathbb{R}^d$ is an unknown parameter vector, and $h^{(t)}(x)$ is a family of carrier densities indexed by $t \in \mathbb{R}$.

The above assumption is a natural extension of a standard exponential family assumption that holds for a single value of t . Specifically, suppose that $h^{(t)}(x)$, $t > 0$, denotes the t -marginal densities of a Lévy process, and that $f_{\theta}^{(1)}(x) = \exp[\theta \cdot x - \psi(\theta)] h^{(1)}(x)$ is an exponential family through $h^{(1)}(x)$ indexed by $\theta \in \mathbb{R}^d$. Then, we can verify that the densities specified in (1.6) induce a family of




Figure 1.4: Illustration of our Poisson process document model with a three-word dictionary and $\mu(\theta) = (0.25, 0.3, 0.45)$. The word counts of the original document, $X = (8, 7, 16)$, represents the trivariate Poisson process A_t , sliced at $T = 28$. The thinned pseudo-document $\tilde{X} = (2, 4, 9)$ represents A_t sliced at $\tilde{T} = 14$.

Lévy processes indexed by θ . The key observation in establishing this result is that, because $h^{(t)}(x)$ is the t -marginal of a Lévy process, the Lévy–Khintchine formula implies that

$$\int e^{\theta \cdot x} h^{(t)}(x) dx = \left(\int e^{\theta \cdot x} h^{(1)}(x) dx \right)^t = e^{t \psi(\theta)},$$

and so the densities in (1.6) are properly normalized.

We also note that, given this assumption and as $T \rightarrow \infty$, we have that A_T/T converges almost surely to $\mu(\theta) \stackrel{\text{def}}{=} \mathbb{E}[A_1]$. Thus, the topic θ can be understood as a description of an infinitely informative input. For finite values of T , X represents a noisy observation of the topic θ .

Now, given this structure, we show that the distribution of $\tilde{X} = A_{\alpha T}$ conditional on $X = A_T$ does not depend on θ . Thus, feature thinning is possible without knowledge of θ using the Lévy thinning procedure defined below. We note that, in our setting, the carrier distributions $h^{(t)}(x)$ are always known; in Section 1.3, we discuss how to efficiently sample from the induced distribution $g^{(\alpha T)}$ for some specific cases of interest.

Theorem 1.1 (Lévy thinning). *Assume that (A_t) satisfies the exponential family structure in (1.6), and let $\alpha \in (0, 1)$ be the thinning parameter. Then, given an input $X = A_T$ and conditioned on any θ , the thinned input $\tilde{X} = A_{\alpha T}$ has the*

following density:

$$g^{(\alpha T)}(\tilde{x}; X) = \frac{h^{(\alpha T)}(\tilde{x}) h^{((1-\alpha)T)}(X - \tilde{x})}{h^{(T)}(X)}, \quad (1.7)$$

which importantly does not depend on θ .

Proof. Because the Lévy process (A_t) has independent and stationary increments, we have that $A_{\alpha T} \sim f_{\theta}^{(\alpha T)}$ and $A_T - A_{\alpha T} \sim f_{\theta}^{((1-\alpha)T)}$ are independent. Therefore, we can write the conditional density of $A_{\alpha T}$ given A_T as the joint density over $(A_{\alpha T}, A_T)$ (equivalently, the reparametrization $(A_{\alpha T}, A_T - A_{\alpha T})$) divided by the marginal density over A_T :

$$\begin{aligned} g^{(\alpha T)}(\tilde{x}; X) &= \frac{f_{\theta}^{(\alpha T)}(\tilde{x}) f_{\theta}^{((1-\alpha)T)}(X - \tilde{x})}{f_{\theta}^{(T)}(X)} \\ &= \left(\exp[\theta \cdot \tilde{x} - \alpha T \psi(\theta)] h^{(\alpha T)}(\tilde{x}) \right) \\ &\quad \times \left(\exp[\theta \cdot (X - \tilde{x}) - (1 - \alpha)T \psi(\theta)] h^{((1-\alpha)T)}(X - \tilde{x}) \right) \\ &\quad \times \left(\exp[\theta \cdot X - T \psi(\theta)] h^{(T)}(X) \right)^{-1}, \end{aligned} \quad (1.8)$$

where the last step expands everything (1.6). Algebraic cancellation, which removes all dependence on θ , completes the proof. \square

Note that while Theorem 1.1 guarantees we can carry out feature thinning without knowing the topic θ , it does not guarantee that we can do it without knowing the information content T . For Poisson processes, the binomial thinning mechanism depends only on α and not on the original T . This is a convenient property in the Poisson case but does not carry over to all Lévy processes — for example, if B_t is a standard Brownian motion, then the distribution of B_2 given $B_4 = 0$ is $\mathcal{N}(0, 1)$, while the distribution of B_{200} given $B_{400} = 0$ is $\mathcal{N}(0, 100)$. As we will see in Section 1.3, thinning in the Gaussian and Gamma families does require knowing T , which will correspond to a “sample size” or “precision.” Likewise, Theorem 1.1 does not guarantee that sampling from (1.7) can be carried out efficiently; however, in all the examples we present here, sampling can be carried out easily in closed form.

1.2.3 Learning with Thinned Features

Having shown how to thin the input X to \tilde{X} without knowledge of θ , we can proceed to defining our full data augmentation strategy. We are given n training examples $\{(X^{(i)}, Y^{(i)})\}_{i=1}^n$. For each original input $X^{(i)}$, we generate B thinned versions $\tilde{X}^{(i,1)}, \dots, \tilde{X}^{(i,B)}$ by sampling from (1.7). We then pair these B examples up with $Y^{(i)}$ and train *any* discriminative classifier on these Bn examples. Algorithm 1 describes the full procedure where we specialize to logistic regression. If one is implementing this procedure using stochastic gradient descent, one can also

Procedure 1. LOGISTIC REGRESSION WITH LÉVY REGULARIZATION

Input: n training examples $(X^{(i)}, Y^{(i)})$, a thinning parameter $\alpha \in (0, 1)$, and a feature map $\phi : \mathbb{R}^d \mapsto \mathbb{R}^p$.

1. For each training example $X^{(i)}$, generate B thinned versions $(\tilde{X}^{(i,b)})_{b=1}^B$ according to (1.7).
2. Train logistic regression on the resulting pseudo-examples:

$$\hat{\beta} \stackrel{\text{def}}{=} \operatorname{argmin}_{\beta \in \mathbb{R}^p \times K} \left\{ \sum_{i=1}^n \sum_{b=1}^B \ell(\beta; \tilde{X}^{(i,b)}, Y^{(i)}) \right\}, \quad (1.9)$$

where the multi-class logistic loss with feature map ϕ is

$$\ell(\beta; x, y) \stackrel{\text{def}}{=} \log \left(\sum_{k=1}^K e^{\beta^{(k)} \cdot \phi(x)} \right) - \beta^{(y)} \cdot \phi(x). \quad (1.10)$$

3. Classify new examples according to

$$\hat{y}(x) = \operatorname{argmin}_{k \in \{1, \dots, K\}} \left\{ \hat{c}^{(k)} - \hat{\beta}^{(k)} \cdot \phi(x) \right\}, \quad (1.11)$$

where the $\hat{c}_k \in \mathbb{R}$ are optional class-specific calibration parameters for $k = 1, \dots, K$.

generate a fresh thinned input \tilde{X} whenever we sample an input X on the fly, which is the usual implementation of dropout training (Srivastava et al., 2014).

In the final step (1.11) of Algorithm 1, we also allow for class-specific calibration parameters. After the $\hat{\beta}^{(k)}$ have been determined by logistic regression with Lévy regularization, these parameters $\hat{c}^{(k)}$ can be chosen by optimizing the logistic loss on the original uncorrupted training data. As discussed in Section 1.2.5, re-calibrating the model is recommended, especially when α is small.

1.2.4 Thinning Preserves the Bayes Decision Boundary

We can easily implement the thinning procedure, but how will it affect the accuracy of the classifier? The following result gives us a first promising piece of the answer by establishing conditions under which thinning does not affect the Bayes decision boundary.

At a high level, our results rely on the fact that under our generative model, the “amount of information” contained in the input vector X is itself uninformative about the class label Y .

Assumption 1.2 (Equal information content across topics). *Assume there exists a constant ψ_0 such that $\psi(\theta) = \psi_0$ with probability 1, over random θ .*

For example, in our Poisson topic model, we imposed the restriction that $\psi(\theta) =$

$\sum_{j=1}^d e^{\theta_j} = 1$, which ensures that the document length $\|A_t\|_1$ has the same distribution (which has expectation $\psi(\theta)$ in this case) for all possible θ .

Theorem 1.2. *Under Assumption 1.2, the posterior class probabilities are invariant under thinning (1.7):*

$$\mathbb{P} [Y = y \mid \tilde{X} = x] = \mathbb{P} [Y = y \mid X = x] \quad (1.12)$$

for all $y \in \{1, \dots, K\}$ and $x \in \mathcal{X}$.

Proof. Given Assumption 1.2, the density of $A_t \mid \theta$ is given by:

$$f_{\theta}^{(t)}(x) = e^{\theta \cdot x} e^{-t\psi_{\theta}} h^{(t)}(x), \quad (1.13)$$

which importantly splits into two factors, one depending on (θ, x) , and the other depending on (t, x) . Now, let us compute the posterior distribution:

$$\mathbb{P} [Y = y \mid A_t = x] \propto \mathbb{P} [Y = y] \int \mathbb{P} [\theta \mid Y] f_{\theta}^{(t)}(x) d\theta \quad (1.14)$$

$$\propto \mathbb{P} [Y = y] \int \mathbb{P} [\theta \mid Y] e^{\theta \cdot x} d\theta, \quad (1.15)$$

which does not depend on t , as $e^{-t\psi_{\theta}} h^{(t)}(x)$ can be folded into the normalization constant. Recall that $X = A_T$ and $\tilde{X} = A_{\tilde{T}}$. Substitute $t = T$ and $t = \tilde{T}$ to conclude (1.12). \square

To see the importance of Assumption 1.2, consider the case where we have two labels ($Y \in \{1, 2\}$), each with a single topic (Y yields topic θ_Y). Suppose that $\psi(\theta_2) = 2\psi(\theta_1)$ —that is, documents in class 2 are on average twice as long as those in class 1. Then, we would be able to make class 2 documents look like class 1 documents by thinning them with $\alpha = 0.5$.

Remark 1.1. *If we also condition on the information content T , then an analogue to Theorem 1.2 holds even without Assumption 1.2:*

$$\mathbb{P} [Y = y \mid \tilde{X} = x, \tilde{T} = t] = \mathbb{P} [Y = y \mid X = x, T = t]. \quad (1.16)$$

This is because, after conditioning on T , the $e^{-t\psi(\theta)}$ term factors out of the likelihood.

The upshot of Theorem 1.2 is that thinning will not induce asymptotic bias whenever an estimator produces $\mathbb{P} [Y = y \mid X = x]$ in the limit of infinite data ($n \rightarrow \infty$), i.e., if the logistic regression (Algorithm 1) is well-specified. Specifically, training either on original examples or thinned examples will both converge to the true class-conditional distribution. The following result assumes that the feature space \mathcal{X} is discrete; the proof can easily be generalized to the case of continuous features.

Corollary 1.3. *Suppose that Assumption 1.2 holds, and that the above multi-class*

logistic regression model is well-specified, i.e., $\mathbb{P}[Y = y | X = x] \propto e^{\beta^{(y)} \cdot \phi(x)}$ for some β and all $y = 1, \dots, K$. Then, assuming that $\mathbb{P}[A_t = x] > 0$ for all $x \in \mathcal{X}$ and $t > 0$, Algorithm 1 is consistent, i.e., the learned classification rule converges to the Bayes classifier as $n \rightarrow \infty$.

Proof. At a fixed x , the population loss $\mathbb{E}[\ell(\beta; X, Y | X = x)]$ is minimized by any choice of β satisfying:

$$\frac{\exp[\beta^{(y)} \cdot \phi(x)]}{\sum_{k=1}^K \exp[\beta^{(k)} \cdot \phi(x)]} = \mathbb{P}[Y = y | X = x] \quad (1.17)$$

for all $y = 1, \dots, K$. Since the model is well-specified and by assumption $\mathbb{P}[\tilde{X} = x] > 0$ for all $x \in \mathcal{X}$, we conclude that weight vector $\hat{\beta}$ learned using Algorithm 1 must satisfy asymptotically (1.17) for all $x \in \mathcal{X}$ as $n \rightarrow \infty$. \square

1.2.5 The End of the Path

As seen above, if we have a correctly specified logistic regression model, then Lévy thinning regularizes it without introducing any bias. However, if the logistic regression model is misspecified, thinning will in general induce bias, and the amount of thinning presents a bias-variance trade-off. The reason for this bias is that although thinning preserves the Bayes decision boundary, it changes the marginal distribution of the covariates X , which in turn affects logistic regression's linear approximation to the decision boundary. Figure 1.5 illustrates this phenomenon in the case where A_t is a Brownian motion, corresponding to Gaussian feature noising; Wager et al. (2014) provides a similar example for the Poisson topic model.

Fully characterizing the bias of Lévy thinning is beyond the scope of this paper. However, we can gain some helpful insights about this bias by studying “strong thinning”—i.e., Lévy thinning in the limit as the thinning parameter $\alpha \rightarrow 0$:

$$\hat{\beta}_{0+} \stackrel{\text{def}}{=} \lim_{\alpha \rightarrow 0} \lim_{B \rightarrow \infty} \hat{\beta}(\alpha, B), \quad (1.18)$$

where $\hat{\beta}(\alpha, B)$ is defined as in (1.9) with the explicit dependence on α and B . For each α , we take $B \rightarrow \infty$ perturbed points for each of the original n data points. As we show in this section, this limiting classifier is well-defined under weak conditions; moreover, in some cases of interest, it can be interpreted as a simple generative classifier. The result below concerns the existence of $\hat{\beta}_{0+}$, and establishes that it is the empirical minimizer of a convex loss function.

Theorem 1.4. *Assume the setting of Procedure 1, and let the feature map be $\phi(x) = x$. Assume that the generative Lévy process (A_t) has finitely many jumps in expectation over the interval $[0, T]$. Then, the limit $\hat{\beta}_{0+}$ is well-defined and can be written as*

$$\hat{\beta}_{0+} = \underset{\beta \in \mathbb{R}^p \times K}{\operatorname{argmin}} \left\{ \sum_{i=1}^n \rho\left(\beta; X^{(i)}, Y^{(i)}\right) \right\}, \quad (1.19)$$




Figure 1.5: The effect of Lévy thinning with data generated from a Gaussian model of the form $X | \theta, T \sim \mathcal{N}(T\theta, \sigma^2 T I_{p \times p})$, as described in Section 1.3.1. The outer circle depicts the distribution of θ conditional on the color Y : blue points all have $\theta \propto (\cos(0.75\pi/2), \sin(0.75\pi/2))$, whereas the red points have $\theta \propto (\cos(\omega\pi/2), \sin(\omega\pi/2))$ where ω is uniform between 0 and $2/3$. Inside this circle, we see 3 clusters of points generated with $T = 0.1, 0.4$, and 1, along with logistic regression decision boundaries obtained from each cluster. The dashed line shows the Bayes decision boundary separating the blue and red points, which is the same for all T (Theorem 1.2). Note that the logistic regression boundaries learned from data with different T are not the same. This issue arises because the Bayes decision boundary is curved, and the best linear approximation to a curved Bayes boundary changes with T .

for some convex function $\rho(\cdot; x, y)$.

The proof of Theorem 1.4 is provided in the appendix. Here, we begin by establishing notation that lets us write down an expression for the limiting loss ρ . First, note that Assumption 1.1 implicitly requires that the process (A_t) has finite moments. Thus, by the Lévy–Itô decomposition, we can uniquely write this process as

$$A_t = bt + W_t + N_t, \quad (1.20)$$

where $b \in \mathbb{R}^p$, W_t is a Wiener process with covariance Σ , and N_t is a compound Poisson process which, by hypothesis, has a finite jump intensity.

Now, by an argument analogous to that in the proof of Theorem 1.1, we see that

the joint distribution of W_T and N_T conditional on A_T does not depend on θ . Thus, we can define the following quantities without ambiguity:

$$\mu_T(x) = bT + \mathbb{E} [W_T \mid A_T = x], \quad (1.21)$$

$$\lambda_T(x) = \mathbb{E} [\text{number of jumps in } (A_t) \text{ for } t \in [0, T] \mid A_T = x], \quad (1.22)$$

$$\nu_T(z; x) = \lim_{t \rightarrow 0} \mathbb{P} [N_t = z \mid N_t \neq 0, A_T = x]. \quad (1.23)$$

More prosaically, $\nu_T(\cdot; x)$ can be described as the distribution of the first jump of \tilde{N}_t , a thinned version of the jump process N_t . In the degenerate case where $\mathbb{P} [\mathbb{N}_T = 0 \mid A_T = x] = 0$, we set $\nu_T(\cdot; x)$ to be a point mass at $z = 0$.

Given this notation, we can write the effective loss function ρ for strong thinning as

$$\begin{aligned} \rho(\beta; x, y) = & -\mu_T(x) \cdot \beta^{(y)} + \frac{T}{2} \frac{1}{K} \sum_{k=1}^K \beta^{(k)\top} \Sigma \beta^{(k)} \\ & + \lambda_T(x) \int \ell(\beta; z, y) d\nu_T(z; x), \end{aligned} \quad (1.24)$$

provided we require without loss of generality that $\sum_{k=1}^K \beta^{(k)} = 0$. In other words, the limiting loss can be described entirely in terms of the distribution of the first jump of \tilde{N}_t , and continuous part W_t of the Lévy process. The reason for this phenomenon is that, in the strong thinning limit, the pseudo-examples $\tilde{X} \sim A_{\alpha T}$ can all be characterized using either 0 or 1 jumps.

Aggregating over all the training examples, we can equivalently write this strong thinning loss as

$$\begin{aligned} \sum_{i=1}^n \rho(\beta; X^{(i)}, Y^{(i)}) = & \frac{1}{2T} \sum_{i=1}^n \gamma_{Y^{(i)}}^{-1} \left\| \gamma_{Y^{(i)}} \mu_T(X^{(i)}) - T \Sigma \beta^{(Y^{(i)})} \right\|_{\Sigma^{-1}}^2 \\ & + \sum_{i=1}^n \lambda_T(X^{(i)}) \int \ell(\beta; z, Y^{(i)}) d\nu_T(z; X^{(i)}), \end{aligned} \quad (1.25)$$

up to $\|\mu_T\|_2^2$ terms that do not depend on β . Here, $\frac{1}{2} \|v\|_{\Sigma^{-1}}^2 = \frac{1}{2} v' \Sigma^{-1} v$ corresponds to the Gaussian log-likelihood with covariance Σ (up to constants), and $\gamma_y = K |\{i : Y^{(i)} = y\}| / n$ measures the over-representation of class y relative to other classes.

In the case where we have the same number of training examples from each class (and so $\gamma_y = 1$ for all $y = 1, \dots, K$), the strong thinning loss can be understood in terms of a generative model. The first term, namely $\frac{1}{2T} \sum_{i=1}^n \left\| \mu_T(X^{(i)}) - T \Sigma \beta^{(Y^{(i)})} \right\|_{\Sigma^{-1}}^2$, is the loss function for linear classification in a Gaussian mixture with observations $\mu_T(X^{(i)})$, while the second term is the logistic loss obtained by classifying single jumps. Thus, strong thinning is effectively seeking the best linear classifier for a generative model that is a mixture of Gaussians and single jumps.

In the pure jump case ($\Sigma = 0$), we also note that strong thinning is closely related

to naive Bayes classification. In fact, if the jump measure of N_t has a finite number of atoms that are all linearly independent, then we can verify that the parameters $\hat{\beta}_{0+}$ learned by strong thinning are equivalent to those learned via naive Bayes, although the calibration constants $c^{(k)}$ may be different.

At a high level, by elucidating the generative model that strong thinning pushes us towards, these results can help us better understand the behavior of Lévy thinning for intermediate value of α , e.g., $\alpha = 1/2$. They also suggest caution with respect to calibration: For both the diffusion and jump terms, we saw above that Lévy thinning gives helpful guidance for the angle of $\beta^{(k)}$, but does not in general elegantly account for signal strength $\|\beta^{(k)}\|_2$ or relative class weights. Thus, we recommend re-calibrating the class decision boundaries obtained by Lévy thinning, as in Algorithm 1.

1.3 Examples

So far, we have developed our theory of Lévy thinning using the Poisson topic model as a motivating example, which corresponds to dropping out words from a document. In this section, we present two models based on other Lévy processes—multivariate Brownian motion (Section 1.3.1) and Gamma processes (Section 1.3.2)—exploring the consequences of Lévy thinning.

1.3.1 Multivariate Brownian Motion

Consider a classification problem where the input vector is the aggregation of multiple noisy, independent measurements of some underlying object. For example, in a biomedical application, we might want to predict a patient's disease status based on a set of biomarkers such as gene expression levels or brain activity. A measurement is typically obtained through a noisy experiment involving an microarray or fMRI, so multiple experiments might be performed and aggregated.

More formally, suppose that patient i has disease status $Y^{(i)}$ and expression level $\mu_i \in \mathbb{R}^d$ for d genes, with the distribution of μ_i different for each disease status. Given μ_i , suppose the t -th measurement for patient i is distributed as

$$Z_{i,t} \sim \mathcal{N}(\mu_i, \Sigma), \quad (1.26)$$

where $\Sigma \in \mathbb{R}^{d \times d}$ is assumed to be a known, fixed matrix. Let the observed input be $X^{(i)} = \sum_{t=1}^{T_i} Z_{i,t}$, the sum of the noisy measurements. If we could take infinitely many measurements ($T_i \rightarrow \infty$), we would have $X^{(i)}/T_i \rightarrow \mu_i$ almost surely; that is, we would observe gene expression noiselessly. For finitely many measurements, $X^{(i)}$ is a noisy proxy for the unobserved μ_i .

We can model the process of accumulating measurements with a multivariate

Brownian motion (A_t):

$$A_t = t\mu + \Sigma^{1/2}B_t, \quad (1.27)$$

where B_t is a d -dimensional white Brownian motion.¹ For integer values of t , A_t represents the sum of the first t measurements, but A_t is also defined for fractional values of t . The distribution of the features X at a given time T is thus

$$X \mid \mu, T \sim \mathcal{N}(T\mu, T\Sigma), \quad (1.28)$$

leading to density

$$\begin{aligned} f_\mu^{(t)}(x) &= \frac{\exp\left[\frac{1}{2}(x - t\mu)^\top (t\Sigma)^{-1}(x - t\mu)\right]}{(2\pi)^{d/2} \det(\Sigma)} \\ &= \exp\left[x^\top \Sigma^{-1}\mu - \frac{t}{2}\mu^\top \Sigma^{-1}\mu\right] h^{(t)}(x), \end{aligned} \quad (1.29)$$

where

$$h^{(t)}(x) = \frac{\exp\left[-\frac{1}{2t}x^\top \Sigma^{-1}x\right]}{(2\pi)^{d/2} \det(\Sigma)^{1/2}}. \quad (1.30)$$

We can recover the form of (1.6) by setting $\theta = \Sigma^{-1}\mu$, a one-to-one mapping provided Σ is positive-definite.

Thinning. The distribution of $\tilde{X} = A_{\alpha T}$ given $X = A_T$ is that of a Brownian bridge process with the following marginals:

$$\tilde{X} \mid X \sim \mathcal{N}(\alpha X, \alpha(1 - \alpha)T\Sigma). \quad (1.31)$$

In this case, “thinning” corresponds exactly to adding zero-mean, additive Gaussian noise to the scaled features αX . Note that in this model, unlike in the Poisson topic model, sampling \tilde{X} from X does require observing T —for example, knowing how many observations were taken. The larger T is, the more noise we need to inject to achieve the same downsampling ratio.

In the Poisson topic model, the features $(X_{i,1}, \dots, X_{i,d})$ were independent of each other given the topic θ_i and expected length T_i . By contrast, in the Brownian motion model the features are correlated (unless Σ is the identity matrix). This serves to illustrate that independence or dependence of the features is irrelevant to our general framework; what is important is that the *increments* $Z_t = A_t - A_{t-1}$ are independent of each other, the key property of a Lévy process.

Assumption 1.2 requires that $\mu^\top \Sigma^{-1} \mu$ is constant across topics; i.e., that the true gene expression levels are equally sized in the Mahalanobis norm defined by Σ . Clearly, this assumption is overly stringent in real situations. Fortunately, Assumption 1.2 is not required (see Remark 1.1) as long as T is observed—as it

1. By definition of Brownian motion, we have marginally that $B_t \sim \mathcal{N}(0, tI)$.

must be anyway if we want to be able to carry out Lévy thinning.

Thinning X in this case is very similar to subsampling. Indeed, for integer values of \tilde{T} , instead of formally carrying out Lévy thinning as detailed above, we could simply resample \tilde{T} values of $Z_{i,t}$ without replacement, and add them together to obtain \tilde{X} . If there are relatively few repeats, however, the resampling scheme can lead to only $\binom{T}{\tilde{T}}$ pseudo-examples (e.g. 6 pseudo-examples if $T = 4$ and $\tilde{T} = 2$), whereas the thinning approach leads to infinitely many possible pseudo-examples we can use to augment the regression. Moreover, if $T = 4$ then subsampling leaves us with only four choices of α ; there would be no way to thin using $\alpha = 0.1$, for instance.

1.3.2 Gamma Process

As another example, suppose again that we are predicting a patient's disease status based on repeated measurements of a biomarker such as gene expression or brain activity. But now, instead of (or in addition to) the *average* signal, we want our features to represent the variance or covariance of the signals across the different measurements.

Assume first that the signals at different genes or brain locations are independent; that is, the t -th measurement for patient i and gene j has distribution

$$Z_{i,j,t} \sim \mathcal{N}(\mu_{i,j}, \sigma_{i,j}^2). \quad (1.32)$$

Here, the variances $\sigma_i^2 = (\sigma_{i,1}^2, \dots, \sigma_{i,d}^2)$ parameterize the “topic.” Suppressing the subscript i , after $T + 1$ measurements we can compute

$$X_{j,T} = \sum_{t=1}^{T+1} (Z_{i,j,t} - \bar{Z}_{i,j,T+1})^2, \quad \text{where} \quad \bar{Z}_{i,j,T+1} = \frac{1}{T+1} \sum_{t=1}^{T+1} Z_{i,j,t}. \quad (1.33)$$

Then $X_{j,T} \sim \sigma_j^2 \chi_T^2$, which is a Gamma distribution with shape parameter $T/2$ and scale parameter $2\sigma_j^2$ (there is no dependence on μ_i). Once again, as we accumulate more and more observations (increasing T), we will have $X_T/T \rightarrow (\sigma_1^2, \dots, \sigma_d^2)$ almost surely.

We can embed $X_{j,T}$ in a multivariate Gamma process with d independent coordinates and scale parameters σ_j^2 :

$$(A_t)_j \sim \text{Gamma}(t/2, 2\sigma_j^2). \quad (1.34)$$

The density of A_t given σ^2 is

$$\begin{aligned} f_{\sigma^2}^{(t)}(x) &= \prod_{j=1}^d \frac{x_j^{t/2-1} e^{-x_j/2\sigma_j^2}}{\Gamma(t/2) 2^{t/2} \sigma_j^{2(t/2)}} \\ &= \exp \left[-\sum_{j=1}^d x_j/2\sigma_j^2 - (t/2) \sum_{j=1}^d \log \sigma_j^2 \right] h^{(t)}(x), \end{aligned} \quad (1.35)$$

where

$$h^{(t)}(x) = \frac{\prod_j x_j^{t/2-1}}{\Gamma(t/2)^d 2^{dt/2}}. \quad (1.36)$$

We can recover the form of (1.6) by setting $\theta_j = -1/2\sigma_j^2$, a one-to-one mapping.

Thinning. Because $\tilde{X}_j \sim \text{Gamma}(\alpha T/2, 2\sigma_j^2)$ is independent of the increment $X_j - \tilde{X}_j \sim \text{Gamma}((1-\alpha)T/2, 2\sigma_j^2)$, we have

$$\frac{\tilde{X}_j}{X_j} \mid X_j \sim \text{Beta}(\alpha T/2, (1-\alpha)T/2). \quad (1.37)$$

In other words, we create a noisy \tilde{X} by generating for each coordinate an independent *multiplicative* noise factor

$$m_j \sim \text{Beta}(\alpha T, (1-\alpha)T) \quad (1.38)$$

and setting $\tilde{X}_j = m_j X_j$. Once again, we can downsample without knowing σ_j^2 , but we do need to observe T . Assumption 1.2 would require that $\prod_j \sigma_j^2$ is identical for all topics. This is an unrealistic assumption, but once again it is unnecessary as long as we observe T .

General covariance. More generally, the signals at different brain locations, or expressions for different genes, will typically be correlated with each other, and these correlations could be important predictors. To model this, let the measurements be distributed as:

$$Z_{i,t} \sim \mathcal{N}(\mu_i, \Sigma_i), \quad (1.39)$$

where Σ represents the unknown “topic”—some covariance matrix that is characteristic of a certain subcategory of a disease status.

After observing $T+1$ observations we can construct the matrix-valued features:

$$X_T = \sum_{t=1}^{T+1} (Z_{i,t} - \bar{Z}_{i,T+1})(Z_{i,t} - \bar{Z}_{i,T+1})^\top. \quad (1.40)$$

Now X_T has a Wishart distribution: $X_T \sim \text{Wish}_d(\Sigma, T)$. When $T \geq d$, the density of A_t given Σ is

$$f_\Sigma^{(t)}(x) = \exp \left\{ -\frac{1}{2} \text{tr}(\Sigma^{-1}x) - \frac{t}{2} \log \det(\Sigma) \right\} h^{(t)}(x), \quad (1.41)$$

where

$$h^{(t)}(x) = \left(2^{\frac{td}{2}} \det(x)^{\frac{t-d-2}{2}} \Gamma_d \left(\frac{t}{2} \right) \right)^{-1}, \quad (1.42)$$

$$\Gamma_d \left(\frac{t}{2} \right) = \pi^{\frac{d(d-1)}{4}} \prod_{j=1}^d \Gamma \left(\frac{t}{2} + \frac{1-j}{2} \right), \quad (1.43)$$

supported on positive-definite symmetric matrices. If $X = A_T$ and $\alpha T \geq d$ as well, we can sample a “thinned” observation \tilde{X} from density proportional to

$$h^{(\alpha T)}(\tilde{x}) h^{(T-\alpha T)}(X - \tilde{x}) \propto \det(\tilde{x})^{\frac{2+d-\alpha T}{2}} \det(X - \tilde{x})^{\frac{2+d-(1-\alpha)T}{2}}, \quad (1.44)$$

or after the affine change of variables $\tilde{X} = X^{1/2} M X^{1/2}$, we sample M from density proportional to $\det(m)^{\frac{2+d-\alpha T}{2}} \det(I_d - m)^{\frac{2+d-(1-\alpha)T}{2}}$, a matrix beta distribution. Here, M may be interpreted as matrix-valued multiplicative noise.

1.4 Simulation Experiments

In this section, we perform several simulations to illustrate the utility of Lévy thinning. In particular, we will highlight the modularity between Lévy thinning (which provides pseudo-examples) and the discriminative learner (which ingests these pseudo-examples). We treat the discriminative learner as a black box, complete with its own internal cross-validation scheme that optimizes accuracy on pseudo-examples. Nonetheless, we show that accuracy on the original examples improves when we train on thinned examples.

More specifically, given a set of training examples $\{(X, Y)\}$, we first use Lévy thinning to generate a set of pseudo-examples $\{(\tilde{X}, Y)\}$. Then we feed these examples to the R function `cv.glmnet` to learn a linear classifier on these pseudo-examples (Friedman et al., 2010). We emphasize that `cv.glmnet` seeks to choose its regularization parameter λ to maximize its accuracy on the pseudo-examples (\tilde{X}, Y) rather than on the original data (X, Y) . Thus, we are using cross-validation as a black box instead of trying to adapt the procedure to the context of Lévy thinning. In principle, we might be concerned that cross-validating on the pseudo-examples would yield a highly suboptimal choice of λ , but our experiments will show that the procedure in fact works quite well.

The two extremes of the path correspond to naive Bayes generative modeling at one end ($\alpha = 0$), and plain ridge-regularized logistic regression at the other ($\alpha = 1$). All methods were calibrated on the training data as follows: Given original weight vectors $\hat{\beta}$, we first compute un-calibrated predictions $\hat{\mu} = X\hat{\beta}$ for the log-odds of $\mathbb{P}[Y = 1 | X]$, and then run a second univariate logistic regression $Y \sim \hat{\mu}$ to adjust both the intercept and the magnitude of the original coefficients. Moreover, when using cross-validation on pseudo-examples (\tilde{X}, Y) , we ensure that all pseudo-examples induced by a given example (X, Y) are in the same cross-validation fold. Code for




Figure 1.6: Performance of Lévy thinning with cross-validated ridge-regularized logistic regression, on a random Gaussian design described in (1.45). The curves depict the relationship between thinning α and classification error as the number of training examples grows: $n = 30, 50, 75, 100, 150, 200, 400$, and 600 . We see that naive Bayes improves over ridge logistic regression in very small samples, while in moderately small samples Lévy thinning does better than either end of the path.

reproducing our results is available at <https://github.com/swager/levythinning>.

Gaussian example. We generate data from the following hierarchical model:

$$Y \sim \text{Binomial}(0.5), \quad \mu | Y \sim \mathcal{L}_Y, \quad X | \mu \sim \mathcal{N}(\mu, I_{d \times d}), \quad (1.45)$$

where $\mu, X \in \mathbb{R}^d$ and $d = 100$. The distribution \mathcal{L}_Y associated with each label Y consists of 10 atoms $\mu_1^{(Y)}, \dots, \mu_{10}^{(Y)}$. These atoms themselves are all randomly generated such that their first 20 coordinates are independent draws of $1.1 T_4$ where T_4 follows Student's t -distribution with 4 degrees of freedom; meanwhile, the last 80 coordinates of μ are all 0. The results in Figure 1.6 are marginalized over the randomness in \mathcal{L}_Y ; i.e., different simulation realizations have different conditional laws for μ given Y . Figure 1.6 shows the results.

Poisson example. We generate data from the following hierarchical model:

$$Y \sim \text{Binomial}(0.5), \quad \theta | Y \sim \mathcal{L}_Y, \quad X_j | \theta \sim \text{Pois} \left(1000 \frac{e^{\theta_j}}{\sum_{j=1}^d e^{\theta_j}} \right), \quad (1.46)$$




Figure 1.7: Performance of Lévy thinning with cross-validated ridge-regularized logistic regression, on a random Poisson design described in (1.46). The curves depict the relationship between thinning α and classification accuracy for $n = 30, 50, 100, 150, 200, 400, 800$, and 1600 . Here, aggressive Lévy thinning with small but non-zero α does substantially better than naive Bayes ($\alpha = 0$) as soon as n is moderately large.

where $\theta \in \mathbb{R}^d$, $X \in \mathbb{N}^d$, and $d = 500$. This time, however, \mathcal{L}_Y is deterministic: If $Y = 0$, then θ is just 7 ones followed by 493 zeros, whereas

$$\theta \mid Y = 1 \sim \left(\underbrace{0, \dots, 0}_7 \mid \underbrace{\tau, \dots, \tau}_7 \mid \underbrace{0, \dots, 0}_{486} \right), \text{ with } \tau \sim \text{Exp}(3).$$

This generative model was also used in simulations by Wager et al. (2014); the difference is that they applied thinning to plain logistic regression, whereas here we verify that Lévy thinning is also helpful when paired with cross-validated ridge logistic regression. Figure 1.7 shows the results.

These experiments suggest that it is reasonable to pair Lévy thinning with a well-tuned black box learner on the pseudo-examples (\tilde{X}, Y) , without worrying about potential interactions between Lévy thinning and the tuning of the discriminative model.

1.5 Discussion

In this chapter, we have explored a general framework for performing data augmentation: apply Lévy thinning and train a discriminative classifier on the resulting pseudo-examples. The exact thinning scheme reflects our generative modeling assumptions. We emphasize that the generative assumptions are non-parametric and of a structural nature; in particular, we never fit an actual generative model, but rather encode the generative hints implicitly in the pseudo-examples.

A key result is that under the generative assumptions, thinning preserves the Bayes decision boundary, which suggests that a well-specified classifier incurs no asymptotic bias. Similarly, we would expect that a misspecified but powerful classifier should incur little bias. We showed that in limit of maximum thinning, the resulting procedure corresponds to fitting a generative model. The exact bias-variance trade-off for moderate levels of thinning is an interesting subject for further study.

While Lévy processes provide a general framework for thinning examples, we recognize that there are many other forms of coarsening that could lead to the same intuitions. For instance, suppose $X \mid \theta$ is a Markov process over words in a document. We might expect that short *contiguous* subsequences of X could serve as good pseudo-examples. More broadly, there are many forms of data augmentation that do not have the intuition of coarsening an input. For example, rotating or shearing an image to generate pseudo-images appeals to other forms of transformational invariance. It would be enlightening to establish a generative framework in which data augmentation with these other forms of invariance arise naturally.

1.6 Appendix: Proof of Theorem 1.4

To establish the desired result, we show that for a single training example (X, Y) , the following limit is well-defined for any $\beta \in \mathbb{R}^{p \times K}$:

$$\begin{aligned} \rho(\beta; X, Y) &= \lim_{\alpha \rightarrow 0} \frac{1}{\alpha} \left(\tilde{\mathbb{E}} \left[\ell(\beta; \tilde{X}, Y) \right] - \log(K) \right) \\ &= -\beta^{(Y)} \cdot X + \lim_{\alpha \rightarrow 0} \frac{1}{\alpha} \tilde{\mathbb{E}} \left[\log \left(\frac{1}{K} \sum_{k=1}^K e^{\beta^{(k)} \cdot \tilde{X}} \right) \right], \end{aligned} \tag{1.47}$$

where on the second line we wrote down the logistic loss explicitly and exploited linearity of the term involving Y as in Wager et al. (2013). Here $\tilde{\mathbb{E}}$ denotes expectation with respect to the thinning process and reflects the $B \rightarrow \infty$ limit. Because ℓ is convex, ρ must also be convex; and by equicontinuity $\hat{\beta}(\alpha)$ must also converge to its minimizer.

Our argument relies on the decomposition $A_t = bt + W_t + N_t$ from (1.20). Without loss of generality, we can generate the pseudo-features \tilde{X} as $\tilde{X} = bt + \tilde{W}_{\alpha T} + \tilde{N}_{\alpha T}$,

where $\widetilde{W}_{\alpha T}$ and $\widetilde{N}_{\alpha T}$ have the same marginal distribution as $W_{\alpha T}$ and $N_{\alpha T}$. Given this notation,

$$\begin{aligned} & \frac{1}{\alpha} \widetilde{\mathbb{E}} \left[\log \left(\frac{1}{K} \sum_{k=1}^K e^{\beta^{(k)} \cdot (\alpha b T + \widetilde{W}_{\alpha T} + \widetilde{N}_{\alpha T})} \right) \right] \\ &= \frac{1}{\alpha} \widetilde{\mathbb{E}} \left[\log \left(\frac{1}{K} \sum_{k=1}^K e^{\beta^{(k)} \cdot (\alpha b T + \widetilde{W}_{\alpha T})} \right) \mid \widetilde{N}_{\alpha T} = 0 \right] \mathbb{P} [\widetilde{N}_{\alpha T} = 0] \\ & \quad + \frac{1}{\alpha} \widetilde{\mathbb{E}} \left[\log \left(\frac{1}{K} \sum_{k=1}^K e^{\beta^{(k)} \cdot (\alpha b T + \widetilde{W}_{\alpha T} + \widetilde{N}_{\alpha T})} \right) \mid \widetilde{N}_{\alpha T} \neq 0 \right] \mathbb{P} [\widetilde{N}_{\alpha T} \neq 0]. \end{aligned}$$

We now characterize these terms individually. First, because N_t has a finite jump intensity, we can verify that, almost surely,

$$\lim_{\alpha \rightarrow 0} \frac{1}{\alpha} \mathbb{P} [\widetilde{N}_{\alpha T} \neq 0] = \lambda_T(X),$$

where $\lambda_T(X)$ is as defined in (1.22). Next, because $\widetilde{W}_{\alpha T}$ concentrates at 0 as $\alpha \rightarrow 0$, we can check that

$$\begin{aligned} & \lim_{\alpha \rightarrow 0} \widetilde{\mathbb{E}} \left[\log \left(\frac{1}{K} \sum_{k=1}^K e^{\beta^{(k)} \cdot (\alpha b T + \widetilde{W}_{\alpha T} + \widetilde{N}_{\alpha T})} \right) \mid \widetilde{N}_{\alpha T} \neq 0 \right] \\ &= \lim_{\alpha \rightarrow 0} \widetilde{\mathbb{E}} \left[\log \left(\frac{1}{K} \sum_{k=1}^K e^{\beta^{(k)} \cdot \widetilde{N}_{\alpha T}} \right) \mid \widetilde{N}_{\alpha T} \neq 0 \right] \\ &= \int \log \left(\frac{1}{K} \sum_{k=1}^K e^{\beta^{(k)} \cdot z} \right) d\nu_T(z; X) \end{aligned}$$

where $\nu_T(\cdot; X)$ (1.23) is the first jump measure conditional on X .

Meanwhile, in order to control the remaining term, we note that we can write

$$\widetilde{W}_{\alpha T} = \alpha \widetilde{W}_T + \widetilde{B}_{\alpha T},$$

where \widetilde{B}_t is a Brownian bridge from 0 to T that is independent from \widetilde{W}_T . Thus, noting that $\lim_{\alpha \rightarrow 0} \mathbb{P} [\widetilde{N}_{\alpha T} = 0] = 1$, we find that

$$\begin{aligned} & \lim_{\alpha \rightarrow 0} \frac{1}{\alpha} \widetilde{\mathbb{E}} \left[\log \left(\frac{1}{K} \sum_{k=1}^K e^{\beta^{(k)} \cdot (\alpha b T + \widetilde{W}_{\alpha T})} \right) \mid \widetilde{N}_{\alpha T} = 0 \right] \mathbb{P} [\widetilde{N}_{\alpha T} = 0] \\ &= \lim_{\alpha \rightarrow 0} \frac{1}{\alpha} \widetilde{\mathbb{E}} \left[\log \left(\frac{1}{K} \sum_{k=1}^K e^{\beta^{(k)} \cdot (\alpha (b T + \widetilde{W}_T) + \widetilde{B}_{\alpha T})} \right) \right] \\ &= \bar{\beta} \cdot \mu_T(X) + \frac{T}{2} \left(\frac{1}{K} \sum_{k=1}^K \beta^{(k) \top} \Sigma \beta^{(k)} - \bar{\beta}^\top \Sigma \bar{\beta} \right), \end{aligned}$$

where $\mu_T(X)$ is as defined in (1.21) and $\bar{\beta} = K^{-1} \sum_{k=1}^K \beta^{(k)}$. The last equality follows from Taylor expanding the $\log(\sum \exp)$ term and noting that 3rd- and higher-

order terms vanish in the limit.

Bringing back the linear term form (1.47), and assuming without loss of generality that $\bar{\beta} = 0$, we finally conclude that

$$\begin{aligned} \rho(\beta; X, Y) &= -\beta^{(Y)} \cdot X + \frac{T}{2} \frac{1}{K} \sum_{k=1}^K \beta^{(k)\top} \Sigma \beta^{(k)} \\ &\quad + \lambda_T(X) \int \log \left(\frac{1}{K} \sum_{k=1}^K e^{\beta^{(k)} \cdot z} \right) d\nu_T(z; X) \\ &= -\beta^{(Y)} \cdot \mu_T(X) + \frac{T}{2} \frac{1}{K} \sum_{k=1}^K \beta^{(k)\top} \Sigma \beta^{(k)} \\ &\quad + \lambda_T(X) \int -\beta^{(Y)} \cdot z + \log \left(\sum_{k=1}^K e^{\beta^{(k)} \cdot z} \right) - \log(K) d\nu_T(z; X), \end{aligned}$$

where for the second equality we used the fact that $X = \mu_T(X) + \lambda_T(X) \int z d\nu_T(z; X)$. Finally, this expression only differs from (1.24) by terms that do not include β ; thus, they yield the same minimizer.

1.7 References

- Y. S. Abu-Mostafa. Learning from hints in neural networks. *Journal of Complexity*, 6(2):192–198, 1990.
- J. Ba and B. Frey. Adaptive dropout for training deep neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, pages 3084–3092, 2013.
- P. Baldi and P. Sadowski. The dropout learning algorithm. *Artificial intelligence*, 210:78–122, 2014.
- C. M. Bishop. Training with noise is equivalent to tikhonov regularization. *Neural computation*, 7(1):108–116, 1995.
- G. Bouchard and B. Triggs. The trade-off between generative and discriminative classifiers. In *International Conference on Computational Statistics*, pages 721–728, 2004.
- B. Efron. The efficiency of logistic regression compared to normal discriminant analysis. *Journal of the American Statistical Association (JASA)*, 70(352):892–898, 1975.
- J. Friedman, T. Hastie, and R. Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1):1–22, 2010.
- A. Globerson and S. Roweis. Nightmare at test time: robust learning by feature deletion. In *International Conference on Machine Learning (ICML)*, pages 353–360, 2006.
- I. Goodfellow, D. Warde-farley, M. Mirza, A. Courville, and Y. Bengio. Maxout networks. In *International Conference on Machine Learning (ICML)*, pages 1319–1327, 2013.
- D. P. Helmbold and P. M. Long. On the inductive bias of dropout. *Journal of*

- Machine Learning Research (JMLR)*, 16:3403–3454, 2015.
- J. Josse and S. Wager. Stable autoencoding: A flexible framework for regularized low-rank matrix estimation. *arXiv preprint arXiv:1410.8275*, 2014.
- A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1097–1105, 2012.
- J. A. Lasserre, C. M. Bishop, and T. P. Minka. Principled hybrids of generative and discriminative models. In *Computer Vision and Pattern Recognition (CVPR)*, pages 87–94, 2006.
- P. Liang and M. I. Jordan. An asymptotic analysis of generative, discriminative, and pseudolikelihood estimators. In *International Conference on Machine Learning (ICML)*, pages 584–591, 2008.
- D. McAllester. A PAC-Bayesian tutorial with a dropout bound. *arXiv preprint arXiv:1307.2118*, 2013.
- A. McCallum, C. Pal, G. Druck, and X. Wang. Multi-conditional learning: Generative/discriminative training for clustering and classification. In *Association for the Advancement of Artificial Intelligence (AAAI)*, 2006.
- A. Y. Ng and M. I. Jordan. On discriminative vs. generative classifiers: A comparison of logistic regression and naive Bayes. In *Advances in Neural Information Processing Systems (NIPS)*, 2002.
- R. Raina, Y. Shen, A. Ng, and A. McCallum. Classification with hybrid generative/discriminative models. In *Advances in Neural Information Processing Systems (NIPS)*, 2004.
- Y. D. Rubinstein and T. Hastie. Discriminative vs informative learning. In *International Conference on Knowledge Discovery and Data Mining (KDD)*, volume 5, pages 49–53, 1997.
- S. P. Schölkopf, P. Simard, V. Vapnik, and A. Smola. Improving the accuracy and speed of support vector machines. In *Advances in Neural Information Processing Systems (NIPS)*, pages 375–381, 1997.
- P. Y. Simard, Y. A. LeCun, J. S. Denker, and B. Victorri. *Transformation Invariance in Pattern Recognition—Tangent Distance and Tangent Propagation*. Neural networks: Tricks of the trade Springer, 1998.
- N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research (JMLR)*, 15(1):1929–1958, 2014.
- L. van der Maaten, M. Chen, S. Tyree, and K. Q. Weinberger. Learning with marginalized corrupted features. In *International Conference on Machine Learning (ICML)*, pages 410–418, 2013.
- S. Wager, S. I. Wang, and P. Liang. Dropout training as adaptive regularization. In *Advances in Neural Information Processing Systems (NIPS)*, 2013.
- S. Wager, W. Fithian, S. I. Wang, and P. Liang. Altitude training: Strong bounds for single-layer dropout. In *Advances in Neural Information Processing Systems (NIPS)*, 2014.
- L. Wan, M. Zeiler, S. Zhang, Y. L. Cun, and R. Fergus. Regularization of neural networks using dropconnect. In *International Conference on Machine Learning (ICML)*, pages 1058–1066, 2013.
- S. I. Wang and C. Manning. Fast dropout training. In *International Conference*

on Machine Learning (ICML), pages 118–126, 2013.

S. I. Wang, M. Wang, S. Wager, P. Liang, and C. Manning. Feature noising for log-linear structured prediction. In *Empirical Methods in Natural Language Processing (EMNLP)*, 2013.